

Robust Overlapping Community Detection via Constrained Egonet Tensor Decomposition

Fatemeh Sheikholeslami
University of Minnesota
Minneapolis, MN
sheik081@umn.edu

Georgios G. Giannakis
University of Minnesota
Minneapolis, MN
georgios@umn.edu

ABSTRACT

The task of community detection over complex networks is of paramount importance in a multitude of applications. The present work puts forward a top-to-bottom community identification approach, termed DC-EgoTen, in which an egonet-tensor (EgoTen) based algorithm is developed in a divide-and-conquer (DC) fashion for breaking the network into smaller subgraphs, out of which the underlying communities progressively emerge. In particular, each step of DC-EgoTen forms a multi-dimensional egonet-based representation of the graph, whose induced structure enables casting the task of overlapping community identification as a constrained PARAFAC decomposition. Thanks to the higher representational capacity of tensors, the novel egonet-based representation improves the quality of detected communities by capturing multi-hop connectivity patterns of the network. In addition, the top-to-bottom approach ensures successive refinement of identified communities, so that the desired resolution is achieved. Synthetic as well as real-world tests corroborate the effectiveness of DC-EgoTen.

This work has been recently published as a regular paper in IEEE Intl. Conf. on Data Mining (ICDM), New Orleans, LA, 2017.

KEYWORDS

Overlapping community detection, egonet subgraphs, tensor decomposition, constrained PARAFAC.

1 INTRODUCTION

Real-world networks often exhibit distinct characteristics, such as power-law degree distribution, the small-world phenomena, and the presence of densely connected sub-graphs, also referred to as “communities” or “clusters” [17]. Focusing on the last, strong connectivity of a subset of nodes along with their sparse interactions with the rest of the network is indicative of a “real-world association” among the participating nodes. The task of community detection targets the discovery of such *communities*, whose identification is of great importance in diverse fields ranging from gene-regulatory networks [12], to brain functionality [33], and social-media evolution analysis [27, 29], to name a few.

Past works on community detection include those based on generative and statistical models [1, 3, 44], modularity and related local-metric optimization [8, 11, 13], spectral clustering [39], and matrix factorization approaches [9, 33, 34, 40, 46, 47]; see also [17] and [15] for comprehensive overviews. However, most existing works pursue a bottom-up approach, where small collections of nodes with strong connectivity patterns (e.g., cliques) are selected as “seeds,” and larger communities are “grown” around them by merging other (clusters of) nodes [11, 42]. In contrast, another class of algorithms follows a top-to-bottom perspective, where a graph is progressively broken into smaller pieces, out of which communities eventually emerge [18, 28, 35].

Recent exploratory studies have revealed new challenges over contemporary networks, addressing the presence of overlapping communities [19, 41, 43], multimodal interaction of nodes over multiview networks [30, 32], exploitation of nodal and edge-related side-information [45], as well as dynamic interactions within a network [4, 5]. In tackling these challenges, *tensors* as multi-modal structures offer increased representational capacity, which translates to improved performance [4, 7, 20, 23, 30, 32, 36].

In this work, we develop a novel top-to-bottom community detection approach, termed “divide-and-conquer EgoTen” (DC-EgoTen), which relies on a successive application of “EgoTen”, a tensor-based toolbox for intermediate steps of community identification. Our core algorithm EgoTen builds on a novel multi-dimensional representation of a network, whose ability in capturing multi-hop connectivities is particularly appealing when communities are overlapping as well as highly-mixing. The proposed tensor-based approach views a network as a union of its egonets, where each egonet is the subgraph induced by a node, its immediate neighbors, and their connections [2]. The resultant three-way tensor is thus built by concatenation of egonet adjacency matrices as frontal slabs. The tensor’s constrained decomposition lends itself to an algorithm revealing communities through the trilinear decomposition factors. A desirable characteristic of this algorithm is its ability to trade off flexibility for increased redundancy and memory costs. Nevertheless, the resulting tensor is extremely sparse, and off-the-shelf tools for sparse tensor computations can be readily utilized; see e.g., [22, 31, 37].

The upshot of our novel framework is three-fold: i) the performance of community detection in complex networks improves markedly thanks to the rich structure of tensors; ii) construction of the egonet-tensor via parallel implementation and exploitation of sparsity endow the algorithm with scalability; and, iii) the proposed top-to-bottom approach offers communities with the desired *resolution*. In fact, many of the previously developed algorithms are susceptible to “resolution limit” [16], where identification of

Work in this paper was supported by NSF grants 1500713, 1442686, 1514056, and NIH grant no. 1R01GM104975-01.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WSDM '18, Heteronam Workshop, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

very large communities reveals little information on the underlying graph structure. The rest of the paper is organized as follows. Sec. 2 introduces DC-EgoTen, and Sec. 3 presents EgoTen as the core tensor-based community detection approach. Extraction of communities and performance metrics are the subjects of Sec. 4, while Sec. 5 presents numerical tests, and Sec. 6 concludes the paper.

2 PRELIMINARIES AND THE TOP-DOWN APPROACH

Given a network of N vertices $v \in \mathcal{V}$ where $|\mathcal{V}| = N$, and their edgeset $\mathcal{E}$, community detection aims at finding subsets of nodes, a.k.a. clusters or communities, for which resident nodes demonstrate dense intra-community connections while distinct communities are sparsely connected. A cover is defined as the set of such communities, with “desirable covers” exhibiting certain characteristics, namely: i) constituent communities should include dense intra-connections and sparse inter-connections; ii) communities of very large sizes are not appealing as they bear little information on the underlying structure of the network; and, iii) the union of the identified communities should cover the entire graph, leaving few or no “homeless” nodes, not assigned to any community.

The proposed “DC-EgoTen” method relies on the construction of an egonet-based multi-dimensional representation of the network. It utilizes “EgoTen” to solve a sequence of nonnegative tensor decomposition subproblems, and progressively unveils the identified communities over the graph. Let us treat EgoTen as a black-box module in this section, postponing its detailed explanation to Section 3, and further delineate the overall algorithm here.

In particular, DC-EgoTen takes a top-down approach for community identification. To this end, “EgoTen” is initially applied over the entire network to provide an assignment of nodes to a few “coarse” communities. Each of the detected communities is in fact a subset of nodes, inducing a subgraph in the overall graph. Thus, the identified “coarse” communities are further amenable to a subsequent application of EgoTen for unraveling a more refined community structure. This procedure can be applied consecutively for a number of times over each of the detected communities, creating a tree of communities, until the desired resolution, i.e., maximum acceptable community size, is achieved for all detected communities (at the leaves of the tree). In Section 3, the proposed egonet-based multi-dimensional graph representation is introduced, and “EgoTen” as our core toolbox for community detection is detailed.

3 EGONET-TENSOR CONSTRUCTION AND CONSTRAINED DECOMPOSITION

Given graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the binary adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is constructed by setting the (i, j) -th entry as $w_{ij} = 1$ if $(i, j) \in \mathcal{E}$, and $w_{ij} = 0$, otherwise. Furthermore, the egonet of node n is defined as the subgraph induced by node n , its one-hop neighbors denoted by $\mathcal{N}(n)$, and all their connections [2]. Thus, the egonet of node n can be conveniently represented by the induced subgraph $\mathcal{G}^{(n)} := (\mathcal{V}, \mathcal{E}^{(n)})$, where $\mathcal{E}^{(n)}$ is the edge set of the links in between nodes $\{n\} \cup \mathcal{N}(n)$. Subsequently, the egonet adjacency matrix $\mathbf{W}^{(n)} \in \mathbb{R}^{N \times N}$ is defined as $w_{ij}^{(n)} := w_{ij}$ if $(i, j) \in \mathcal{E}^{(n)}$, and $(i, j) \in \mathcal{E}^{(n)} = 0$




Figure 1: Construction of the three-way egonet-tensor.

Algorithm 1 Egonet-tensor construction

```

procedure EGONET-TENSOR CONSTRUCTION( $\mathcal{V}, \mathbf{W}$ )
  for  $n \in \mathcal{V}$  do
     $\mathcal{N}(n) := \{v \in \mathcal{V} | w_{nv} \neq 0\}$ 
     $\mathbf{W}^n \leftarrow \text{subgraph}(\{n\} \cup \mathcal{N}(n), \mathbf{W})$ 
     $\underline{\mathbf{W}}_{:,n} = \mathbf{W}^{(n)}$ 
  end for
end procedure
return  $\underline{\mathbf{W}}$ 

```

otherwise. Typically, the center node n is excluded from $\mathcal{G}^{(n)}$, but it is included here for convenience.

Let us now consider a *three-way egonet-tensor* $\underline{\mathbf{W}} \in \mathbb{R}^{N \times N \times N}$ constructed by concatenating egonet adjacency matrices $\mathbf{W}^{(n)}$ for all nodes $n \in \mathcal{V}$ in the frontal slabs of $\underline{\mathbf{W}}$. In tensor parlance, that is tantamount to setting the n -th frontal slab of $\underline{\mathbf{W}}$ as $\underline{\mathbf{W}}_{:,n} := \mathbf{W}^{(n)}$, where $:$ is a free index that spans its range.

The advantage of representing a graph via its egonet-tensor is due to the fact that tensors as multi-way data structures are capable of capturing higher-order connectivities, namely two-hop links among neighboring nodes. Thus, in networks where overlapping as well as highly-mixed communities render the task of community detection very challenging, egonet-tensors provide a rich representation of the graph, which will be leveraged in the upcoming algorithm. The egonet-based representation is also of interest particularly in the absence of extra nodal features, as the enhanced representation is a result of careful exploitation of the adjacency matrix where no other source of information is provided.

Fig. 1 illustrates the egonet-tensor construction procedure, while Alg. 1 provides its pseudocode. In the ensuing subsection we cast the task of community detection as a constrained tensor decomposition over the egonet-tensor $\underline{\mathbf{W}}$, elaborate on the intuition behind the proposed approach, and introduce *EgoTen* as its efficient solver.

3.1 EgoTen: A Constrained Tensor Decomposition Approach

In order to gain insights into the properties of the introduced egonet-tensor, consider the toy network whose connectivity is depicted in Figure 2a. The network under consideration comprises five communities with dense intra-community and fewer inter-community connections. Upon constructing the egonet-tensor and after permutation (so that resident nodes are indexed right after one another), it becomes evident that the egonet-tensor demonstrates a block structure; see Fig. 2b. In particular, dense diagonal blocks in the tensor capture the dense intra-community links, while sparse off-diagonal entries represent inter-community connections.




Figure 2: (a) A toy network with 5 non-overlapping communities; (b) corresponding egonet-tensor; and (c) its community-revealing factorization via PARAFAC.

Had the communities been complete sub-graphs, each block would have been an all-one three-way tensor (considering non-zero diagonal entries provided by self-loops), which could have been readily decomposed into the outer product of three all-one vectors (each of the size of the community); that is, $\mathbf{1}_{p \times p \times p} = \mathbf{1}_{p \times 1} \circ \mathbf{1}_{p \times 1} \circ \mathbf{1}_{p \times 1}$, where p is the size of the community. Moreover, had the communities been disjoint, that is if no inter-community links were present, the egonet-tensor could have been readily written as the summation of five tensors, each of whom can be effectively approximated by the outer-product of three vectors; see Fig. 2c.

Such decomposition is indeed reminiscent of the well-known canonical polyadic decomposition (CPD) [22] also known as PARAFAC, where the number of terms, i.e., the rank of the decomposition, reveals the number of communities. Prompted by this observation, let us introduce the constrained nonnegative PARAFAC over the egonet adjacency tensor $\underline{\mathbf{W}}$ as

$$\begin{aligned} \{\hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}\} = & \arg \min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} \left\{ \|\underline{\mathbf{W}} - \sum_{k=1}^K \mathbf{a}_k \circ \mathbf{b}_k \circ \mathbf{c}_k\|_F^2 \right. \\ & \left. + \lambda (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2) \right\} \\ \text{s.t. } & \mathbf{A} \geq \mathbf{0}, \mathbf{B} \geq \mathbf{0}, \mathbf{C} \geq \mathbf{0} \\ & \sum_{k=1}^K c_{nk} = 1 \quad \forall n = 1, 2, \dots, N \end{aligned} \quad (1)$$

The first term in the objective is the original Frobenious term in the well-known PARAFAC, through which minimization of the mismatch between the multi-way data structure $\underline{\mathbf{W}}$ and its approximation is achieved. Furthermore, nonnegativity of the egonet-tensor is effected through additional constraints over the factors $\mathbf{A} := [\mathbf{a}_1, \dots, \mathbf{a}_K]$, $\mathbf{B} := [\mathbf{b}_1, \dots, \mathbf{b}_K]$ and $\mathbf{C} := [\mathbf{c}_1, \dots, \mathbf{c}_K]$. Regarding the simplex constraints on the rows of matrix $\mathbf{C}$, let us now focus on the n -th frontal slab of the egonet-tensor. One can readily show that the tensor approximation gives rise to following decomposition

$$\mathbf{W}^{(n)} \simeq \sum_{k=1}^K c_{nk} (\mathbf{a}_k \circ \mathbf{b}_k) \quad (2)$$

where c_{nk} denotes the (n, k) -th entry of factor $\mathbf{C}$. As stated earlier, parameter K is referred to as the rank of the decomposition, and in

this application reveals the number of identified communities. Thus, such decomposition can be interpreted as a weighted sum over K “basis” $\{\mathbf{a}_k \circ \mathbf{b}_k\}_{k=1}^K$, where $(\mathbf{a}_k \circ \mathbf{b}_k)$ captures the “connectivity structure” within the k -th community. Consequently, c_{nk} can be viewed as *association level* of node n to community k . Thus, the simplex constraint over the rows of matrix $\mathbf{C}$ readily guarantees a normalized association vector for every node in the graph to the identified K communities. Finally, the Frobenious regularizers over factors $\mathbf{A}$ and $\mathbf{B}$ simply resolve the scaling ambiguity between the two factors, and is different from [6]. The overall optimization in (1) is a block-convex problem [21], whose solver is discussed next.

3.2 Constrained PARAFAC Solver

Exploiting the block-convex structure of the constrained PARAFAC in (1), the optimization can be solved by alternating minimization, where each of $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ is optimized respectively by fixing the other two at their current values. Factors are repeatedly updated until a stopping criterion or a maximum number of iterations is achieved. Considering iteration i , factors are updated as follows.

3.2.1 Factor A update. Fixing factors $\mathbf{B}^{(i-1)}$ and $\mathbf{C}^{(i-1)}$ at their current values, the update of factor $\mathbf{A}$ is obtained by the corresponding subproblem, which after algebraic manipulation can be readily rewritten as a regularized nonnegative least-squares (LS) minimization as

$$\mathbf{A}^{(i)} = \arg \min_{\mathbf{A} \geq \mathbf{0}} \|\mathbf{W}_1 - \mathbf{H}_A^{(i)} \mathbf{A}^\top\|_F^2 + \lambda \|\mathbf{A}\|_F^2 \quad (3)$$

where $\mathbf{W}_1 := [\text{vec}(\underline{\mathbf{W}}_{1,:}), \dots, \text{vec}(\underline{\mathbf{W}}_{N,:})] \in \mathbb{R}^{N^2 \times N}$ is a matrix-cized reshaping of the tensor $\underline{\mathbf{W}}$, and $\mathbf{H}_A^{(i)} := [\mathbf{b}_1^{(i-1)} \otimes \mathbf{c}_1^{(i-1)}, \dots, \mathbf{b}_K^{(i-1)} \otimes \mathbf{c}_K^{(i-1)}]$, with $\mathbf{b}_c^{(i-1)}$ ($\mathbf{c}_c^{(i-1)}$) denoting column c of $\mathbf{B}^{(i-1)}$ (resp. $\mathbf{C}^{(i-1)}$), and $\otimes$ the Kronecker product operator; see also [22]. Solving the subproblem in (3) by the alternating direction method of multipliers (ADMM), the updates are obtained as

$$\begin{cases} \bar{\mathbf{A}}^{(r)} = \arg \min_{\bar{\mathbf{A}}} \mathcal{L}_A^{(i)}(\mathbf{A}^{(r-1)}, \bar{\mathbf{A}}, \mathbf{Y}^{(r-1)}) \\ = \left(\mathbf{H}_A^{(i)\top} \mathbf{H}_A^{(i)} + (\lambda + \rho/2) \mathbf{I}_{K \times K} \right)^{-1} \\ \quad \times \left(\mathbf{W}_1^\top \mathbf{H}_A^{(i)} + \frac{\rho}{2} (\mathbf{Y}^{(r-1)} + \mathbf{A}^{(r-1)}) \right) \\ \mathbf{A}^{(r)} = \mathcal{P}_+(\mathbf{Y}^{(r-1)} - \bar{\mathbf{A}}^{(r)}) \\ \mathbf{Y}^{(r)} = \mathbf{Y}^{(r-1)} - \rho(\mathbf{A}^{(r)} - \bar{\mathbf{A}}^{(r)}) \\ r = r + 1 \end{cases} \quad (4)$$

where $\bar{\mathbf{A}}, \mathbf{Y} \in \mathbb{R}^{N \times K}$ are the auxiliary and dual variables, respectively, and regularization parameter is set as $\rho = \|\mathbf{H}_A^{(i)}\|_F^2 / K$. Also, $\mathcal{P}_+(\cdot)$ denotes the element-wise projection onto the nonnegative subspace. The iterations proceed until $\|\mathbf{A}^{(r)} - \mathbf{A}^{(r-1)}\| / \|\mathbf{A}^{(r-1)}\| \leq \epsilon$, or the maximum number of iterations is exceeded. Upon its termination, factor $\mathbf{A}$ is updated as $\mathbf{A}^{(i)} \leftarrow \mathbf{A}^{(r)}$, and the algorithm proceeds with updating factor $\mathbf{B}$ as in the following.

3.2.2 Factor B update. Upon fixing $\mathbf{A} = \mathbf{A}^{(i)}$ and $\mathbf{C} = \mathbf{C}^{(i-1)}$, factor $\mathbf{B}$ is updated by solving the subproblem

$$\mathbf{B}^{(i)} = \arg \min_{\mathbf{B} \geq \mathbf{0}} \|\mathbf{W}_2 - \mathbf{H}_B^{(i)} \mathbf{B}^\top\|_F^2 + \lambda \|\mathbf{B}\|_F^2 \quad (5)$$

where $\mathbf{W}_2 := [\text{vec}(\mathbf{W}_{:,1,:}), \dots, \text{vec}(\mathbf{W}_{:,N,:})] \in \mathbb{R}^{N^2 \times N}$, and $\mathbf{H}_B^{(i)} := [\mathbf{a}_1^{(i)} \otimes \mathbf{c}_1^{(i-1)}, \dots, \mathbf{a}_K^{(i)} \otimes \mathbf{c}_K^{(i-1)}]$, yielding a similar optimization problem as in (3). Undertaking the same approach as for (4), the ADMM update for solving (5) yields

$$\begin{cases} \bar{\mathbf{B}}^{(r)} = \left(\mathbf{H}_B^{(i)\top} \mathbf{H}_B^{(i)} + (\lambda + \rho/2) \mathbf{I}_{K \times K} \right)^{-1} \\ \quad \times \left(\mathbf{W}_2^\top \mathbf{H}_B^{(i)} + \frac{\rho}{2} (\mathbf{Y}^{(r-1)} + \mathbf{B}^{(r-1)}) \right) \\ \mathbf{B}^{(r)} = \mathcal{P}_+ \left(\mathbf{Y}^{(r-1)} - \bar{\mathbf{B}}^{(r)} \right) \\ \mathbf{Y}^{(r)} = \mathbf{Y}^{(r-1)} - \rho (\mathbf{B}^{(r)} - \bar{\mathbf{B}}^{(r)}) \\ r = r + 1. \end{cases} \quad (6)$$

Upon the termination of (6) due to either attaining the stopping criterion or reaching the maximum number of iterations, factor $\mathbf{B}$ is updated as $\mathbf{B}^{(i)} \leftarrow \mathbf{B}^{(r)}$.

3.2.3 Factor C update. Fixing factors $\mathbf{A} = \mathbf{A}^{(i)}$ and $\mathbf{B} = \mathbf{B}^{(i)}$, update of factor $\mathbf{C}$ is obtained by solving the subproblem

$$\begin{aligned} \mathbf{C}^{(i)} &= \arg \min_{\mathbf{C}} \|\mathbf{W}_3 - \mathbf{H}_C^{(i)} \mathbf{C}^\top\|_F^2 \\ \text{s.t. } &\mathbf{C} \geq \mathbf{0} \quad \sum_{k=1}^K c_{nk} = 1 \quad \forall n = 1, \dots, N \end{aligned} \quad (7)$$

where $\mathbf{W}_3 := [\text{vec}(\mathbf{W}_{:,1,:}), \dots, \text{vec}(\mathbf{W}_{:,N,:})]$ is the matricized version of $\mathbf{W}$ along the 3-rd mode, and $\mathbf{H}_C^{(i)} := [\mathbf{a}_1^{(i)} \otimes \mathbf{b}_1^{(i)}, \dots, \mathbf{a}_K^{(i)} \otimes \mathbf{b}_K^{(i)}]$. ADMM solver proceeds with the following iterative updates

$$\begin{cases} \bar{\mathbf{C}}^{(r)} = \arg \min_{\mathbf{C}} \mathcal{L}_C^{(i)}(\mathbf{C}^{(r-1)}, \bar{\mathbf{C}}, \mathbf{Y}^{(r-1)}) \\ \quad = \left(\mathbf{H}_C^{(i)\top} \mathbf{H}_C^{(i)} + \rho/2 \mathbf{I}_{K \times K} \right)^{-1} \\ \quad \times \left(\mathbf{W}_3^\top \mathbf{H}_C^{(i)} + \frac{\rho}{2} (\mathbf{Y}^{(r-1)} + \mathbf{C}^{(r-1)}) \right) \\ \mathbf{C}^{(r)} = \mathcal{P}_{\text{simp}}(\mathbf{Y}^{(r-1)} - \bar{\mathbf{C}}^{(r)}) \\ \mathbf{Y}^{(r)} = \mathbf{Y}^{(r-1)} - \rho (\mathbf{C}^{(r)} - \bar{\mathbf{C}}^{(r)}) \\ r = r + 1. \end{cases} \quad (8)$$

Projection of the rows of matrix $(\mathbf{Y}^{(r-1)} - \bar{\mathbf{C}}^{(r)})$ onto the simplex set can be achieved via the algorithm in [14]. Upon termination, factor $\mathbf{C}$ is updated as $\mathbf{C}^{(i)} \leftarrow \mathbf{C}^{(r)}$.

Once the overall trilinear optimization in (1) is solved, factor $\mathbf{C}$ unravels soft community association of the nodes. Extraction of hard communities based on the learned PARAFAC model is discussed in the next section. Also, Algorithm 2 lists the pseudocode of the proposed EgoTen followed by hard community assignments.

4 COMMUNITY ASSIGNMENT

As discussed in Section 3, the introduced EgoTen community detection algorithm aims at solving a constrained decomposition of the egonet-tensor, thus providing factor $\mathbf{C}$ whose entries unravel soft community associations. In order to transform the “soft” to “hard” memberships, one can simply utilize a threshold approach, according to which if $c_{nk} > \tau_k$, node n is assigned to community k , and it is not assigned otherwise. The main challenge here is on selecting a proper threshold τ_k . To this end, let $\hat{\mathcal{C}}_k$ denote the set of nodes in community k (with hard memberships), and define its conductance as [17]

$$\phi(\hat{\mathcal{C}}_k) := \frac{\sum_{i \in \hat{\mathcal{C}}_k, j \notin \hat{\mathcal{C}}_k} \mathbf{W}_{ij}}{\min\{\text{vol}(\hat{\mathcal{C}}_k), \text{vol}(\mathcal{V} \setminus \hat{\mathcal{C}}_k)\}}$$

Algorithm 2 EgoTen Community Detection Core Algorithm

```

procedure EgoTen( $\mathbf{W}, K$ )
  Initialize  $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{R}^{N \times K}$  at random and set  $i = 0$ 
  while  $i < I_{\max}$  do or not-converged
     $\mathbf{A}^{(i)} \leftarrow$  Solve (3) using (4)
     $\mathbf{B}^{(i)} \leftarrow$  Solve (5) using (6)
     $\mathbf{C}^{(i)} \leftarrow$  Solve (7) using (8)
     $i \leftarrow i + 1$ 
  end while
  for  $k = 1, 2, \dots, K$  do
     $\hat{\mathcal{C}}_k = \{\}$ 
    for  $n = 1, 2, \dots, N$  do  $\hat{\mathcal{C}}_k \leftarrow \hat{\mathcal{C}}_k \cup \{n\}$  if  $c_{nk} \geq \tau_k$ 
    end for
  end for
end procedure
retrun  $\{\hat{\mathcal{C}}_k\}_{k=1}^K$ 

```

where $\text{vol}(\hat{\mathcal{C}}_k) := \sum_{i \in \hat{\mathcal{C}}_k, \forall j} \mathbf{W}_{ij}$ and $(\mathcal{V} \setminus \hat{\mathcal{C}}_k)$ is the complement of $\hat{\mathcal{C}}_k$. According to $\phi(\cdot)$, *high-quality* communities yield small conductance scores as they exhibit dense connections among the nodes within the community and sparse connections with the rest.

Considering conductance as a measure of community quality, we can now set threshold τ_k such that the quality of community k after hard member assignment is maximized. In order to lower complexity, we simply choose τ_k from the discretized range $[1/K, 2/K, \dots]$. Note that having an association level $c_{nk} = 1/K \forall k$ for a given node n is tantamount to having an equally favorable association with the K communities, and having threshold $\tau_k = 1/K$ will result in a community assignment if the association is higher than this uniform level. Also, setting $\tau_k = 1/K$ together with the simplex constraints on the rows of factor $\mathbf{C}$ guarantees that every node will be assigned to at least one community, and no node will be left unassigned. However, tuning τ_k to obtain low conductance communities improves quality.

4.1 DC-EgoTen

Having delineated different modules of DC-EgoTen, we are ready to present the overall algorithm. Given graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, DC-EgoTen initially constructs the egonet-tensor $\mathbf{W}$ using Alg. 1, applies EgoTen in Alg. 2 over $\mathbf{W}$, and obtains detected communities $\{\hat{\mathcal{C}}_k\}_{k=1}^K$. Next, the resolution of $\hat{\mathcal{C}}_k$ for $k = 1, 2, \dots$ will determine whether further refining is necessary for each of the identified communities. That is, if $|\hat{\mathcal{C}}_k| < C_{\max}$, the resolution of detected community $\hat{\mathcal{C}}_k$ is satisfactory, and no further processing is required. On the other hand, if $|\hat{\mathcal{C}}_k| > C_{\max}$, the subgraph induced by the set of nodes in $\hat{\mathcal{C}}_k$ will be extracted, over which the entire process will be repeated. Alg/ 3 lists the pseudocode for the overall DC-EgoTen.

Figure 3 provides a schematic over our toy network with five communities, each of size $|\mathcal{C}_k| = 15$ for $k = 1, 2, \dots, 5$. In this example, in every EgoTen the rank parameter is $K = 2$, which gives rise to a binary tree of detected communities. As in this example, in the first application of EgoTen, the green community is detected by the constrained PARAFAC, while the rest of the network is ‘lumped’ together in the second community. Thus, the green community




Figure 3: The proposed DC-EgoTen on a toy example.

Algorithm 3 DC-EgoTen

```

procedure DC-EgoTen( $\mathcal{V}, \mathbf{W}$ )
    Set parameters  $K, C_{\max}$ 
    Define global cover set  $\mathcal{S} = \{\}$ 
     $\mathbf{W} \leftarrow \text{Egonet-tensor construction}(\mathcal{V}, \mathbf{W})$ 
     $\{C_i\}_{i=1,2,\dots,K} \leftarrow \text{EgoTen}(\mathbf{W}, K)$ 
    for  $C \in \{C_i\}_{i=1,2,\dots,K}$  do
        # If community  $C$  is refined enough, add it to the cover
        set  $\mathcal{S}$ , otherwise refine it using EgoTen
        if  $|C| < C_{\max}$  then
             $\mathcal{S} \leftarrow \mathcal{S} \cup C$ 
        else
            # Extract the subgraph of nodes in  $C$ 
             $\mathbf{W}_{\text{sub}} \leftarrow \text{subgraph}(C, \mathbf{W})$ 
            DC-EgoTen( $C, \mathbf{W}_{\text{sub}}$ )
        end if
    end for
end procedure
return  $\mathcal{S}$ 

```

needs no further processing as its size is below $C_{\max} = 20$, while application of EgoTen on the second term gives rise to two relatively more refined communities. Proceeding with another set of EgoTen application on the detected communities will reveal the remaining clusters, creating overall five leaves in the tree, corresponding to the detected fine-resolution communities.

If an oracle had provided the number of underlying communities, the algorithm would have identified all clusters in its first application of EgoTen by setting $K = 5$. However, successive application of EgoTen with smaller target rank K can compensate for the lack of such information, which is almost-always encountered in practice. Furthermore, DC-EgoTen nicely proceeds with the desiderata of community identification discussed in Section 2, because: i) the multi-dimensional egonet-based representation captures multi-hop connectivities, leading to an improved quality in the detected communities; ii) consecutive division of large communities enhances resolution; and, iii) setting threshold parameter $\tau_k = 1/K$ in EgoTen can guarantee a full coverage of the network, while its tuning can further control the trade-off between coverage and quality.

4.2 Performance Evaluation

In addition to conductance, normalized mutual information and F1-score are measures for assessing the performance of community identification when ground-truth communities are provided.

Normalized mutual information (NMI) [15]: Given $\mathcal{S}^* = \{C_1^*, \dots, C_{|\mathcal{S}^*|}^*\}$ and $\hat{\mathcal{S}} = \{\hat{C}_1, \dots, \hat{C}_{|\hat{\mathcal{S}}|}\}$ as ground-truth and detected covers, respectively, the information theoretic measure NMI is defined as (cf. [15])

$$\text{NMI}(\mathcal{S}^*, \hat{\mathcal{S}}) := \frac{2I(\mathcal{S}^*, \hat{\mathcal{S}})}{H(\mathcal{S}^*) + H(\hat{\mathcal{S}})}$$

where $H(\hat{\mathcal{S}})$ denotes the entropy of set $\hat{\mathcal{S}}$ defined as

$$H(\hat{\mathcal{S}}) := - \sum_{i=1}^{|\hat{\mathcal{S}}|} p(\hat{C}_i) \log p(\hat{C}_i) = - \sum_{i=1}^{|\hat{\mathcal{S}}|} \frac{|\hat{C}_i|}{N} \log \frac{|\hat{C}_i|}{N}$$

and similarly for $H(\mathcal{S}^*)$. Furthermore, $I(\mathcal{S}^*, \hat{\mathcal{S}})$ denotes the mutual information between $\mathcal{S}^*$ and $\hat{\mathcal{S}}$, defined as

$$I(\mathcal{S}^*, \hat{\mathcal{S}}) := \sum_{i=1}^{|\mathcal{S}^*|} \sum_{j=1}^{|\hat{\mathcal{S}}|} \frac{|C_i^* \cap \hat{C}_j|}{N} \log \frac{N|C_i^* \cap \hat{C}_j|}{|C_i^*||\hat{C}_j|}. \quad (9)$$

Intuitively, the mutual information $I(\mathcal{S}^*, \hat{\mathcal{S}})$ reflects a measure of similarity between the two covers. Thus, high values of NMI, namely its maximum at 1, reflect high accuracy in community identification, whereas low values of NMI, namely its minimum at 0, represent poor discovery of the true underlying communities. This measure has been generalized for overlapping communities in [24], and will be utilized for performance assessment in such networks.

Average F1-score [44]: F1-score is a measure of binary classification accuracy. Specifically, the harmonic mean of *precision* and *recall* takes its highest value at 1 and lowest value at 0. Average F1-score for detected cover $\hat{\mathcal{S}}$ is

$$\bar{F1} := \frac{1}{2|\mathcal{S}^*|} \sum_{i=1}^{|\mathcal{S}^*|} F1(C_i^*, \hat{C}_{I(i)}) + \frac{1}{2|\hat{\mathcal{S}}|} \sum_{i=1}^{|\hat{\mathcal{S}}|} F1(C_{I'(i)}^*, \hat{C}_i)$$

where $I(i) = \arg \max_j F1(C_i^*, \hat{C}_j)$, $I'(i) = \arg \max_j F1(C_j^*, \hat{C}_i)$ in

$$\text{which } F1(C_i, C_j) := \frac{2|C_i \cap C_j|}{|C_i| + |C_j|}.$$

5 NUMERICAL TESTS

In this section, the proposed DC-EgoTen is applied to synthetic as well as real datasets. Synthetic Lancichinetti-Fortunato and Radicci (LFR) networks [25] are utilized as a benchmark to study the resilience and performance of different community identification algorithms with overlapping as well as mixing communities.

5.1 LFR Benchmark Networks

LFR graphs serve as benchmark networks in which certain real-world properties, namely power-law distribution for nodal degree and community sizes, as well as the presence of overlapping and mixing communities are preserved. Such networks are configured by a total number of N nodes, $\bar{d}$ average degree, and power-law distribution exponents γ_1 and γ_2 for degree and community sizes, respectively. Furthermore, parameter μ controls the community mixing, where higher values result in more out-of-community edges in

between non-resident nodes. Moreover, parameters o_n, o_m respectively set the number of overlapping nodes and communities (with which these nodes are associated).

In order to assess the resilience of the proposed DC-EgoTen to variations of μ and o_n , we have generated networks with $N = 2,000$, $\bar{d} = 100$, $\gamma_1 = 2$, $\gamma_2 = 1$, and varied $\mu \in [0.1, 0.7]$ as well as o_n in 10% – 70% of the total networks size N , respectively. DC-EgoTen is run by setting the rank K in the initial application as $K = 100$, while following applications are set as $K = 2$, essentially leading to a bisection of the network in the subsequent steps, and sparse tensor decompositions are handled via the SPLATT toolbox [38]. Thresholding parameter τ_k is selected as explained in Section 4 for the top EgoTen (allowing for overlapping community detection), and set as $\tau = 1/2$ for next steps. Maximum community size is set as $C_{\max} = 200$. The performance is compared with state-of-the-art algorithms BigClam [44], Demon [10], and Nise [42] with ‘spread-hub’ seeding strategy, where $|\hat{S}| = 200$ is provided as an estimate on the number of communities in Nise and BigClam. Due to the availability of underlying communities, the performance is assessed via NMI and F1-scores and averaged over 10 realizations of the network for each setting. As the results in Figures 4 and 5 corroborate, DC-EgoTen provides higher performance in terms of NMI and F1-score, thanks to the rich egonet-based representation as well as the progressive identification of refined communities.

5.2 Real-world Networks

In this subsection, the performance of DC-EgoTen is compared with state-of-the-art overlapping community detection algorithms on various real-world networks, listed in Table I, available in [26]. In DC-EgoTen, constructing the egonet-tensors as well as solving the constrained PARAFAC utilize parallel implementation, while BigClam and Nise also allow for parallel threading. Thus, for networks with $N < 1$ million, these algorithms are run using 8 threads and 32GB of RAM, while for the Youtube dataset, 24 threads with 256 GB of RAM are utilized. As with synthetic datasets, we apply DC-EgoTen with $K = 100$ for the first application of EgoTen, and set $K = 2$ for subsequent steps. Threshold τ_k is selected as explained in Sec. 4 for the top EgoTen (allowing for overlapping community detection), and $\tau = 1/2$ for next steps. Also, maximum community size $C_{\max}$ is set to 1% of the network size for each dataset.

Figure 6 plots the run time of different algorithms while Table II lists the coverage and number of detected communities. Due to unavailability of ground-truth communities, NMI and F1-score could not be evaluated, thus performance is assessed using the




Figure 4: Performance of different algorithms versus different community mixing values μ for $o_n = 600$, and $o_m = 3$.




Figure 5: Performance of different algorithms versus different number of overlapping nodes o_n for $\mu = 0.2$, and $o_m = 3$.

Table 1: Real-world networks.

Dataset	No. of vertices N	No. of edges $ \mathcal{E} $	Edge type
Facebook	4,039	88,234	Undirected
Enron	36,692	183,831	Undirected
Epinion	75,879	508,837	Directed
Slashdot	82,168	948,464	Directed
Email	265,214	420,045	Directed
Stanford	281,903	2,312,497	Directed
Notredame	325,729	1,497,134	Directed
Youtube	1,134,890	2,987,624	Undirected

conductance-coverage curve. To this end, for a given algorithm, the conductance of the identified communities is computed and the communities are sorted accordingly in an increasing order. Conductance-coverage curve is then plotted by increasing the maximum conductance, and progressively adding the sorted communities to the set of covered nodes. Figure 7 depicts the aforementioned curve for various datasets. As low values of conductance

Table 2: Coverage and number of detected communities of different methods over real-world networks.

Dataset		DC-EgoTen	Bigclam	Demon	Nise
Facebook	Coverage	100%	95%	99%	89%
	# comm.	523	500	8	16
Enron	Coverage	100%	90%	65%	100%
	# comm.	553	500	343	520
Slashdot	Coverage	100%	100%	95%	100%
	# comm.	1163	500	51	485
Epinion	Coverage	100%	100%	35%	100%
	# comm.	1274	2000	136	2041
Email	Coverage	100 %	83%	11%	100%
	# comm.	965	2000	24	2404
Notredame	Coverage	100%	100%	39%	100%
	# comm.	1169	2000	1497	1454
Stanford	Coverage	100%	90%	85%	100%
	# comm.	807	2000	2596	1411
Youtube	Coverage	100%	100%	22%	100%
	# comm.	813	5000	3835	5162

correspond to more cohesive communities, a smaller area under curve (AUC) generally implies better performance. However, the resolution of the communities is another important metric which must be considered in drawing conclusions. Interestingly, the separation of different scattered points for a given algorithm in the conductance-coverage curve reveals the granularity of the detected communities. That is, if a detected community is very large, its inclusion creates a jump in the coverage, which is noticeable by the two consecutive points in the plot being placed far apart. Thus, examining Figure 7 reveals that the identified communities via DC-EgoTen and Bigclam are usually of more refined sizes as those plots are always smooth, while the performance of Nise and Demon is often limited to detecting very large communities (upto 40% of the whole network). Furthermore, although one may not particularly be interested in 100% coverage, it is desirable that a relatively high number of nodes to be covered within the detected communities, and thus low coverage where more than 50% of the nodes are left uncovered is considered undesirable.

6 CONCLUSION

This work dealt with identification of overlapping communities via DC-EgoTen, a top-to-bottom tensor-based framework. Specifically, a novel egonet-based tensor representation of a network was introduced and utilized in a constrained PARAFAC decomposition, whose factors subsequently reveal the underlying communities. To provide the detected communities with desirable resolution, this algorithm was applied progressively in a top-to-bottom fashion, where the network is decomposed into K communities per step. Parallel implementation as well as exploitation of the sparsity in the egonet-tensor endow the algorithm with scalability, while the structured redundancy and the rich representational capacity of the egonet-tensor enhance the performance of the toolbox. Sparse sampling of egonets along the third mode is among our future directions, through which memory as well as computational requirements of the algorithm can be reduced, while the structured redundancy in the egonet-tensor is expected to preserve performance.




Figure 6: Runtime of different algorithms on various datasets denoted on the x-axis as: (D1) Facebook, (D2) Enron, (D3) Epinion, (D4) Email, (D5) Slashdot, (D6) Notredame, (D7) Stanford, and (D8) Youtube.

REFERENCES

- [1] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Vancouver, Canada, Dec. 2009. Mixed membership stochastic blockmodels. *Proc. of Advances in Neural Inf. Proc. Systems* (Vancouver, Canada, Dec. 2009), 33–40.
- [2] L. Akoglu, M. McGlohon, and C. Faloutsos. Washington, DC, USA, 2010. Oddball: Spotting anomalies in weighted graphs. *Advances in Knowledge Discovery and Data Mining* (Washington, DC, USA, 2010), 410–421.
- [3] A. Anandkumar, R. Ge, D. Hsu, and S. M. Kakade. Jan. 2014. A tensor approach to learning mixed membership community models. *J. of Machine Learning Research* 15, 1 (Jan. 2014), 2239–2312.
- [4] Miguel Araujo, Spiros Papadimitriou, Stephan Günnemann, Christos Faloutsos, Prithwish Basu, Ananthram Swami, Evangelos E Papalexakis, and Danai Koutra. Tairan Taiwan, Aug. 2014. Com2: Fast automatic discovery of temporal (āY-cometāZ) communities. *Pacific-Asia Conf. on Knowledge Discovery and Data Mining* (Tairan Taiwan, Aug. 2014), 271–283.
- [5] Brian Baingana and Georgios B Giannakis. 2016. Joint community and anomaly tracking in dynamic networks. *IEEE Trans. on Signal Proc.* 64, 8 (2016), 2013–2025.
- [6] J. A. Bazerque, G. Mateos, and G. B. Giannakis. Nov. 2013. Rank Regularization and Bayesian Inference for Tensor Completion and Extrapolation. *IEEE Trans. on Signal Proc.* 61, 22 (Nov. 2013), 5689–5703.
- [7] A. R. Benson, D. F. Gleich, and J. Leskovec. Vancouver, Canada, Feb. 2015. Tensor spectral clustering for partitioning higher-order network structures. *Proc. of SIAM Intl. Conf. on Data Mining*, (Vancouver, Canada, Feb. 2015), 118–126.
- [8] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008:1–12.
- [9] Xiaochun Cao, Xiao Wang, Di Jin, Yixin Cao, and Dongxiao He. 2013. Identifying overlapping communities as well as hubs and outliers via nonnegative matrix factorization. *Scientific Reports* 3 (2013), 2993.
- [10] Michele Coscia, Giulio Rossetti, Fosca Giannotti, and Dino Pedreschi. 2012. Demon: A local-first discovery method for overlapping communities. (2012), 615–623.
- [11] I. Derényi, G. Palla, and T. Vicsek. 2005. Clique percolation in random networks. *Physical Review Letters* 94, 16 (2005), 160202–1:4.
- [12] Pan Du, Jian Gong, E. S. Wurtel, and J. A. Dickerson. 2005. Modeling gene expression networks using fuzzy logic. *IEEE Trans. on Systems, Man, and Cybernetics, Part B* 35, 6 (Dec 2005), 1351–1359. <https://doi.org/10.1109/TSMCB.2005.855590>
- [13] Jordi Duch and Alex Arenas. 2005. Community detection in complex networks using extremal optimization. *Physical Review E* 72, 2 (2005), 027104:1–4.
- [14] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. 2008. Efficient projections onto the l1-ball for learning in high dimensions. *Proc. Intl. Conf. on Machine Learning* (2008), 272–279.
- [15] S. Fortunato. Feb. 2010. Community detection in graphs. *Physics Reports* 486, 3 (Feb. 2010), 75–174.
- [16] Santo Fortunato and Marc Barthelemy. 2007. Resolution limit in community detection. *Proc. of the National Academy of Sciences* 104, 1 (2007), 36–41.
- [17] Santo Fortunato and Darko Hric. 2016. Community detection in networks: A user guide. *Physics Reports* 659 (2016), 1–44.
- [18] Steve Gregory. Warsaw, Poland, 2007. An algorithm to find overlapping community structure in networks. *European Conf. on Principles of Data Mining and Knowledge Discovery* (Warsaw, Poland, 2007), 91–102.
- [19] Kun He, Yiwel Sun, David Bindel, John Hopcroft, and Yixuan Li. 2015. Detecting overlapping communities from local spectral subspaces. *IEEE Intl. Conf. on Data*




Figure 7: Conductance-coverage curve for various datasets using different community detection algorithms.

- Mining (2015), 769–774.
- [20] Fulong Huang, UN Niranjan, Mohammad Umar Hakeem, and Animashree Anandkumar. 2015. Online tensor methods for learning latent variable models. *J. of Machine Learning Research* 16 (2015), 2797–2835.
- [21] K. Huang, N. D. Sidiropoulos, and A. P. Liavas. 2016. A Flexible and Efficient Algorithmic Framework for Constrained Matrix and Tensor Factorization. *IEEE Trans. on Signal Processing* 64, 19 (2016), 5052–5065.
- [22] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM Rev.* 51, 3 (2009), 455–500.
- [23] Tamara G Kolda, Brett W Bader, and Joseph P Kenny. Houston, Texas, 2005. Higher-order web link analysis using multilinear algebra. *IEEE Intl. Conf. on Data Mining* (Houston, Texas, 2005), 8–pp.
- [24] Andrea Lancichinetti, Santo Fortunato, and János Kertész. 2009. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics* 11, 3 (2009), 033015.
- [25] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. 2008. Benchmark graphs for testing community detection algorithms. *Physical Review E* 78, 4 (2008), 046110.
- [26] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>. (June 2014).
- [27] Y. R. Lin, Y. Chi, S. Zhu, H. Sundaram, and B. L. Tseng. Paris, France, April 2009. Analyzing communities and their evolutions in dynamic social networks. *ACM Trans. on Knowledge Discovery from Data* 3, 2 (Paris, France, April 2009), 8.
- [28] Mark E J Newman. 2004. Fast algorithm for detecting community structure in networks. *Physical Review E* 69, 6 (2004).
- [29] Symeon Papadopoulos, Yiannis Kompatsiaris, Athena Vakali, and Ploutarchos Spyridonos. Beijing, China, 2012. Community detection in social media. *Data Mining and Knowledge Discovery* 24, 3 (Beijing, China, 2012), 515–554.
- [30] Evangelos E Papalexakis, Leman Akoglu, and Dino Ience. Istanbul, Turkey, 2013. Do more views of a graph help? Community detection and clustering in multi-graphs. *IEEE Intl. Conf. on Information Fusion* (Istanbul, Turkey, 2013), 899–905.
- [31] Evangelos E Papalexakis, Christos Faloutsos, and Nicholas D Sidiropoulos. 2012. Parcube: Sparse parallelizable tensor decompositions. *Joint European Conf. on Machine Learning and Knowledge Discovery in Databases* (2012), 521–536.
- [32] E. E. Papalexakis, N. D. Sidiropoulos, and R. Bro. Jan. 2013. From K-Means to Higher-Way Co-Clustering: Multilinear Decomposition With Sparse Latent Factors. *IEEE Trans. on Signal Proc.* 61, 2 (Jan. 2013), 493–506.
- [33] J. D. Power, A. L. Cohen, S. M. Nelson, G. S. Wig, K. A. Barnes, J. A. Church, A. C. Vogel, T. O. Laumann, F. M. Miezin, B. L. Schlaggar, and S. E. Petersen. 2011. Functional network organization of the human brain. *Neuron* 72, 4 (2011), 665–678.
- [34] Ioannis Psorakis, Stephen Roberts, Mark Ebden, and Ben Sheldon. 2011. Overlapping community detection using Bayesian non-negative matrix factorization. *Physical Review E* 83, 6 (2011), 066114.
- [35] Matthew J Rattigan, Marc Maier, and David Jensen. Corvallis, OR, USA, 2007. Graph clustering with network structure indices. In *Proc. Intl. Conf. on Machine Learning*. 783–790.
- [36] Fatemeh Sheikholeslami, Brian Baingana, Georgios B. Giannakis, and Nikolaos D. Sidiropoulos. DC, Dec. 2016. Egonet tensor decomposition for community identification. *Proc. of Globasis* (DC, Dec. 2016).
- [37] Shaden Smith, Alec Beri, and George Karypis. Bristol, UK, Aug. 2017. Constrained Tensor Factorization with Accelerated AO-ADMM. In *IEEE Intl. Conf. on Parallel Processing*.
- [38] Shaden Smith and George Karypis. [n. d.]. SPLATT: The Surprisingly Parallel Sparse Tensor Toolkit. <http://cs.umn.edu/~splatt/>. ([n. d.]).
- [39] U. Von Luxburg. Aug. 2007. A tutorial on spectral clustering. *Statistics and Computing* 17, 4 (Aug. 2007), 395–416.
- [40] F. Wang, T. Li, X. Wang, S. Zhu, and C. Ding. Paris France, July 2011. Community discovery using nonnegative matrix factorization. *ACM Trans. on Data Mining and Knowledge Discovery* 22, 3 (Paris France, July 2011), 493–521.
- [41] Joyce Jiyoung Whang, Inderjit S Dhillon, and David F Gleich. 2015. Non-exhaustive, overlapping k-means. *Proc. of SIAM Intl. Conf. on Data Mining* (2015), 936–944.
- [42] Joyce Jiyoung Whang, David F. Gleich, and Inderjit S. Dhillon. 2016. Overlapping community detection using neighborhood-inflated seed expansion. *IEEE Trans. on Knowledge and Data Engineering* 28, 5 (May 2016), 1272–1284.
- [43] Joyce Jiyoung Whang, David F. Gleich, and Inderjit S. Dhillon. San Diego, CA, USA, 2013. Overlapping community detection using seed set expansion. *Proc. of ACM Intl. Conf. on Information and Knowledge Management* (San Diego, CA, USA, 2013), 2099–2108.
- [44] J. Yang and J. Leskovec. Rome, Italy, Feb. 2013. Overlapping community detection at scale: A nonnegative matrix factorization approach. *Proc. of ACM Intl. Conf. on Web Search and Data Mining* (Rome, Italy, Feb. 2013), 587–596.
- [45] Jaewon Yang, Julian McAuley, and Jure Leskovec. Dallas, TX, USA, 2013. Community detection in networks with node attributes. *IEEE Intl. Conf. on Data Mining* (Dallas, TX, USA, 2013), 1151–1156.
- [46] Yu Zhang and Dit Yan Yeung. Beijing, China, 2012. Overlapping community detection via bounded nonnegative matrix tri-factorization. *Proc. of ACM Intl. Conf. on Knowledge Discovery and Data Mining* (Beijing, China, 2012), 606–614.
- [47] Zhong Yuan Zhang, Yong Wang, and Yong Yeol Ahn. 2013. Overlapping community detection in complex networks using symmetric binary matrix factorization. *Physical Review E* 87, 6 (2013), 062803.